arXiv:1301.4055v1 [math.CO] 17 Jan 2013

Heat-bath Markov chains have no negative eigenvalues

Catherine Greenhill*
School of Mathematics and Statistics
The University of New South Wales
Sydney NSW 2052, Australia
`csg@unsw.edu.au`

Mario Ullrich†
Dipartimento di Matematica
Università Roma Tre
00146 Rome, Italy
`ullrich.mario@gmail.com`

January 18, 2013

Abstract

We prove a result which says that when the transition matrix of a Markov chain has a certain structure, then the chain has no negative eigenvalues. For such chains, it is not necessary to make the chain lazy in order to investigate the mixing time. The lemma is designed to apply to heat-bath Markov chains. We present two applications of this result: one to single-site heat-bath chains for spin systems and one to a heat-bath Markov chain for sampling contingency tables.

1 The result and two applications

Suppose that Ω is a finite set and let $\pi : \Omega \to (0, 1]$ be a probability distribution on Ω. Let $\mathcal{L}$ be a nonempty finite index set and let $L = |\mathcal{L}|$. Suppose that for all $x \in \Omega$ and $a \in \mathcal{L}$ we have a subset $\Omega_{x,a}$ of Ω such that

- $x \in \Omega_{x,a}$ for all $x \in \Omega$ and $a \in \mathcal{L}$, and
- for each $a \in \mathcal{L}$, the set $\{\Omega_{x,a} : x \in \Omega\}$ forms a partition of Ω.

These conditions imply that for all $x, y \in \Omega$ and $a \in \mathcal{L}$,

$$\text{if} \quad y \in \Omega_{x,a} \quad \text{then} \quad \Omega_{x,a} = \Omega_{y,a}. \tag{1}$$

Now suppose that $\mathcal{M}$ is a Markov chain on Ω with transition matrix P given by

$$P = \frac{1}{L} \sum_{a \in \mathcal{L}} P_a \tag{2}$$

---

*Research supported by the Australian Research Council.

†Research supported by ERC Advanced Grant PTRELSS.

where, for each $a \in \mathcal{L}$, the $|\Omega| \times |\Omega|$ matrix P_a (with rows and columns indexed by Ω) is defined by

$$P_a(x,y) = \frac{\pi(y)}{\pi(\Omega_{x,a})}\, \mathbb{1}(y \in \Omega_{x,a}).$$

(Here $\mathbb{1}(y \in \Omega_{x,a})$ equals 1 if $y \in \Omega_{x,a}$, and equals 0 otherwise.) Note that P_a is well-defined for all $a \in \mathcal{L}$, since π is nonzero on all states and $\Omega_{x,a}$ is nonempty.

If $x \in \Omega$ is the current state of the Markov chain $\mathcal{M}$ then a transition is performed by choosing an element $a \in \mathcal{L}$ uniformly at random, then sampling the next state y from $\Omega_{x,a}$ with respect to the distribution π restricted to $\Omega_{x,a}$. (To make the connection with heat-bath chains, think of $\mathcal{L}$ as the set of positions where updates may occur, and for a given $x \in \Omega$ and $a \in \mathcal{L}$, the set $\Omega_{x,a}$ contains the possibilities for the next state.)

When (2) holds it follows that $\mathcal{M}$ is aperiodic (since every state has a self-loop) and that $\mathcal{M}$ is reversible with respect to π. However, the chain $\mathcal{M}$ need not be irreducible. (See [7] for Markov chain definitions which are not given here.)

Lemma 1. *Let $\mathcal{M}$ be a Markov chain with finite state space Ω and transition matrix P. Let $\pi : \Omega \to (0,1]$ be a probability distribution on Ω. Suppose that P satisfies (2) for some index set $\mathcal{L}$, sets $\Omega_{x,a}$ and matrices P_a. Then $\mathcal{M}$ has no negative eigenvalues.*

Proof. Recall that a matrix is *positive semi-definite* if it has no negative eigenvalues. By definition of P_a we know that $P_a(x,y) = 0$ if $y \notin \Omega_{x,a}$. Furthermore (1) implies that if $z \in \Omega_{x,a}$ then $P_a(x,y) = P_a(z,y)$ for all $y \in \Omega$. Therefore, for all $x, y \in \Omega$ and all $a \in \mathcal{L}$ we have

$$\begin{aligned} P_a^2(x,y) = \sum_{z\in\Omega} P_a(x,z)\,P_a(z,y) = \sum_{z\in\Omega_{x,a}} P_a(x,z)\,P_a(z,y) &= \sum_{z\in\Omega_{x,a}} P_a(x,z)\,P_a(x,y) \\ &= P_a(x,y). \end{aligned}$$

This says that the matrix P_a is a projection, so the only eigenvalues of P_a are 0 and 1. It follows that P_a is positive semi-definite for all $a \in \mathcal{L}$.

Then P is a nonnegative linear combination of positive semi-definite matrices, and it follows that P is positive semi-definite. (This can be proved directly from the variational characterisation of eigenvalues, or see for example [6, Observation 7.1.3].) This completes the proof. $\square$

Lemma 1 can be easily generalised to allow a non-uniform distribution over $\mathcal{L}$, with the same proof. That is, we can replace (2) by

$$P = \sum_{a\in\mathcal{L}} \rho(a)\, P_a,$$

where ρ is a probability distribution on $\mathcal{L}$.

1.1 Application: a single-site heat-bath chain for spin systems

Let $G = (V, E)$ be an arbitrary graph and let S be a finite set of spins (or colors). Consider the state space $\Omega = S^V$ and let π be a probability distribution on Ω. Given $\sigma \in \Omega$, for all $v \in V$ and $k \in S$ we define $\sigma^{v,k}$ by

$$\sigma^{v,k}(u) = \begin{cases} \sigma(u) & \text{if } u \neq v, \\ k & \text{otherwise.} \end{cases}$$

(So $\sigma^{v,k}$ is obtained from σ by replacing the spin at v by k.) The single-site *heat-bath chain* for Ω is the Markov chain with transition matrix

$$P(\sigma, \sigma^{v,k}) = \frac{1}{|V|} \sum_{v \in V} \frac{\pi(\sigma^{v,k})}{\sum_{\ell \in S} \pi(\sigma^{v,\ell})}.$$

This matches the setting of Lemma 1 by choosing $\mathcal{L} = V$ and $L = |V|$, and defining

$$\Omega_{\sigma,v} = \{\sigma^{v,k} : k \in S\}$$

for all $\sigma \in \Omega$ and $v \in V$. Hence, by Lemma 1, single-site heat-bath chains for general spin models do not have negative eigenvalues. $\diamond$

The heat-bath chain, which belongs to the family of Glauber dynamics (see for example [9]), has been studied by many authors including [5, 8, 10]. In several cases, the continuous-time version of this Markov chain is considered. One advantage of this approach is that mixing properties can be described solely by the second-largest eigenvalue. Thus, when translating these results to discrete time, it usually remains to bound the smallest eigenvalue of the chain. The last example shows that for the heat-bath chain, the established continuous-time bounds can be used without further analysis. In the case of the Potts model this argument was used in the proof of [12, Theorem 2.10].

Note that there are other Glauber dynamics, like the Metropolis chain, which are generally not guaranteed to have only non-negative eigenvalues.

1.2 Application: a heat-bath chain for contingency tables

Let $\mathbf{r} = (r_1, \dots, r_m)$ and $\mathbf{c} = (c_1, \dots, c_n)$ be two vectors of positive integers with the same sum. A *contingency table* with row sums r and column sums c is an $m \times n$ matrix with nonnegative integer entries, such that the i'th row sum is r_i and the j'th column sum is c_j, for $i = 1, \dots, m$ and $j = 1, \dots, n$. Let $\Omega_{\mathbf{r},\mathbf{c}}$ denote the set of all contingency tables with row sums $\mathbf{r}$ and column sums $\mathbf{c}$.

Dyer and Greenhill [4] proposed a Markov chain for sampling contingency tables, which we will call the *contingency chain*. A transition of the chain is performed as follows: choose a 2×2 subsquare of the current table uniformly at random, then replace this 2×2 subsquare by a uniformly chosen 2×2 nonnegative integer matrix with the same row and column sums. The *lazy* contingency chain does nothing at each step with probability $\frac{1}{2}$, and otherwise performs a transition as described above. Cryan et al. [1] analysed the lazy contingency chain for a constant number of rows and proved that it is rapidly mixing.

To fit the contingency chain into the setting of Lemma 1, let $\mathcal{L}$ be the set of all positions of 2×2 subsquares, and let $L = |\mathcal{L}| = \binom{m}{2}\binom{n}{2}$. Let P_a be the transition matrix of the Markov chain which acts only on the 2×2 subsquare $a \in \mathcal{L}$. Then $P_a(x, \cdot)$ is uniform over all contingency tables $y \in \Omega_{\mathbf{r},\mathbf{c}}$ which differ from x only within the 2×2 subsquare a. Hence Lemma 1 applies and shows that the contingency chain has no negative eigenvalues.

2 Implications for the mixing time

Let $\mathcal{M}$ be an ergodic, reversible Markov chain with finite state space Ω, transition matrix P and stationary distribution π. The eigenvalues of $\mathcal{M}$ satisfy

$$1 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{N-1} > -1,$$

where $N = |\Omega|$. We refer to λ_{N-1} as the *smallest eigenvalue* of $\mathcal{M}$. (Note, this does not mean smallest in absolute value.) The connection between the mixing time of a Markov chain and its eigenvalues is well-known (see [11, Proposition 1]):

$$\tau(\varepsilon) \leq (1 - \lambda_*)^{-1} \ln \frac{1}{\epsilon\, \pi_{\min}} \tag{3}$$

where $\tau(\varepsilon)$ denotes the mixing time of the Markov chain, $\pi_{\min} = \min_{x \in \Omega} \pi(x)$ and

$$\lambda_* = \max\{\lambda_1,\, |\lambda_{N-1}|\}.$$

When studying the mixing time of a Markov chain $\mathcal{M}$, the approach which has become standard is to make the chain $\mathcal{M}$ *lazy* by replacing P by $(I + P)/2$, where I denotes the identity matrix. Then all eigenvalues of the lazy chain are nonnegative, and only the second-largest eigenvalue must be investigated. Clearly if P has no negative eigenvalues then $\lambda_* = \lambda_1$ and it is not necessary to make the chain lazy.

Diaconis and Stroock [3, Proposition 2] gave a result which can be used to calculate an upper bound for $(1 + \lambda_{N-1})^{-1}$ for ergodic, reversible chains. A modified version was stated by Diaconis and Saloffe-Coste [2, p.702]. In particular, if every state has a self-loop then

$$(1 + \lambda_{N-1})^{-1} \leq \tfrac{1}{2} \max_{x \in \Omega} P(x, x)^{-1}.$$

For many combinatorial Markov chains, this approach is simple to apply and gives an upper bound on $(1 + \lambda_{N-1})^{-1}$ which is several orders of magnitude smaller than the best-known bound on $(1 - \lambda_1)^{-1}$. We feel that these results are under-utilised, and recommend them as an alternative to making Markov chains lazy (together with Lemma 1, where it applies).

References

[1] M. Cryan, M. Dyer, L.A. Goldberg, M. Jerrum and R. Martin, Rapidly mixing Markov chains for sampling contingency tables with a constant number of rows, *SIAM Journal on Computing* **36** (2006), 247–278.

[2] Diaconis and L. Saloffe-Coste, Comparison theorems for reversible Markov chains, *Annals of Applied Probability* **3** (1993), 696–730.

[3] P. Diaconis and D. Stroock, Geometric bounds for eigenvalues of Markov chains, *Annals of Applied Probability* **1** (1991), 36–61.

[4] M.E. Dyer and C. Greenhill, Polynomial-time counting and sampling of two-rowed contingency tables, *Theoretical Computer Science* **246** (2000), 265–278.

[5] T. Hayes, A simple condition implying rapid mixing of single-site dynamics on spin systems, in *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, IEEE, Los Alimitos, 2006, pp. 39–46.

[6] R.A. Horn and C.R. Johnson, *Matrix Analysis*, Cambridge University Press, Cambridge, 1985.

[7] M. Jerrum, *Counting, Sampling and Integrating: algorithms and complexity*, Lectures in Mathematics – ETH Zürich, Birkhäuser, Basel, 2003.

[8] E. Lubetzky and A. Sly. Critical Ising on the square lattice mixes in polynomial time. *Comm. Math. Phys.*, 2012. doi:10.1007/s00220-012-1460-9.

[9] F. Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on probability theory and statistics (Saint-Flour, 1997)*, volume 1717 of *Lecture Notes in Math.*, pages 93–191. Springer, Berlin, 1999.

[10] F. Martinelli and E. Olivieri, Approach to equilibrium of Glauber dynamics in the one phase region I. The attractive case, *Communications in Mathematical Physics* **161** (1994), pp. 447–486.

[11] A. Sinclair, Improved bounds for mixing rates of Markov chains and multicommodity flow, *Combinatorics, Probability and Computing* **1** (1992), 351–370.

[12] M. Ullrich, Rapid mixing of Swendsen-Wang dynamics in two dimensions, *PhD thesis, Friedrich Schiller University, Jena* (2012).